arXiv:1604.06752v1 [cs.LO] 22 Apr 2016

Solving QBF by Abstraction*

Leander Tentrup

Reactive Systems Group
Saarland University

Abstract. Many synthesis and verification approaches rely on solving techniques for quantified Boolean formulas (QBF). Consequently, solution witnesses, i.e., Boolean functions, become more important as they represent implementations or counterexamples. We present a new recursive counterexample guided abstraction refinement algorithm for solving and certifying QBFs that exploits structural reasoning on the formula level. The algorithm decomposes the given QBF into one propositional formula for every scope—a maximal block of consecutive quantifiers of the same type—that abstracts from assignments of outer and inner quantifiers. Further, we define a format for proof traces and present an efficient method to extract certificates from traces that allows building witnessing functions and certifying the solver result. We report on experimental evaluation of this algorithm in a prototype solver QuAbS (Quantified Abstraction Solver) using synthesis benchmark sets and show the effectiveness of the certification approach.

1 Introduction

Efficient solving techniques for Boolean theories are an integral part of modern verification and synthesis methods. The ever growing complexity of verification and synthesis problems led to propositional problems of enormous size. To see further advances in these areas, we believe that it is necessary to move to more compact representations of these requirements. Quantified Boolean formulas (QBFs) have repeatedly been considered as a candidate theory to compactly encode Boolean problems [3, 9, 10, 13, 19, 29, 32, 38, 42]. Recent advances in QBF solvers give rise to hope that QBF may help to increase the scalability of verification and synthesis approaches.

Most current QBF solving algorithms are limited to formulas in prenex conjunctive normal form (PCNF). Starting with formulas in CNF enables efficient optimizations and preprocessing methods, like blocked clause elimination [8], that are more difficult to implement on more general structures. However, there is experimental evidence that CNF solvers can be ineffective even when the satisfiability of the formula is trivial [1, 41]. Further, problems are rarely naturally

* This work was partially supported by the German Research Foundation (DFG) as part of the Transregional Collaborative Research Center "Automatic Verification and Analysis of Complex Systems" (SFB/TR 14 AVACS).

expressed in CNF and the CNF transformation increases the formula's size linearly, disrupts its structure, and artificially introduces new variables. We argue that negation normal form (NNF) is a good compromise between structure and complexity. Arbitrary quantified Boolean formulas can be transformed to NNF with maximal increase of formula size of factor two while preserving most of the structure and without artificially adding variables. At the same time, NNF inherits important properties from CNF, especially that negations appear only in front of variables, enabling pure literal analysis and other optimizations that we describe further in Sec. 3.2.

Certifying the results is particularly important for QBF, as the pure yes/no answer is of little use. Like the propositional SAT problem [7, 12], QBF can be used to encode objects of interest, like error paths [13,29] and implementations [9, 19, 20]. While the yes/no answer of a SAT or QBF solver then provides the information about the *existence* of this object, we often want to construct a concrete instance for further use. For the propositional SAT problem the object can typically be extracted as the assignment of the variables, but for QBF the object potentially consists of *Skolem* functions for the existential quantifiers or *Herbrand* functions for the universal quantifiers. Certification has been explored in the setting of CNF solving [2,24,34,37] but in practice it involves performance penalties due to non-applicable solving optimizations and limited preprocessing. We believe that certification must be treated as a first-class citizen of QBF solving and show that it is possible to have competitive performance and solution extraction at the same time.

We present a new counterexample guided abstraction refinement (CEGAR) algorithm for solving quantified Boolean formulas that exploits the propositional formula's structure. The algorithm decomposes the given QBF into one propositional formula for every scope—a maximal block of consecutive quantifiers of the same type—that *abstracts* from assignments of outer and inner scopes. This propositional abstraction contains special interface literals that are used to communicate abstraction entries between the scopes. In every scope, we use a SAT solver as an oracle to generate new abstraction entries and to provide us with witnesses for unsatisfiable queries. On the algorithmic side, we introduce a new element to QBF refinement algorithms by maintaining, for every scope, a *dual abstraction* and make twofold use of it: It provides a method for optimizing abstraction entries and it is used to translate counterexamples from inner scopes to outer scopes.

To sum up, this paper makes the following contributions:

- We provide a general counterexample guided abstraction refinement (CEGAR) algorithm for solving QBFs in negation normal form.
- We provide requirements that an abstraction must fulfill to be compatible with this algorithm and provide an instance of an abstraction that is experimentally evaluated.
- We show an exponential gap between the abstraction algorithm and algorithms based on assignment and expansion.

– We describe an efficient certification approach, i.e., proof traces and function extraction. It is evaluated on synthesis benchmarks where implementations can be directly obtained from certificates.

2 Quantified Boolean Formulas

A quantified Boolean formula (QBF) is a propositional formula over a finite set of variables $\mathcal{X}$ with domain $\mathbb{B} = \{0, 1\}$ extended with quantification. The syntax is given by the grammar

$$\varphi \coloneqq x \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \exists x.\, \varphi \mid \forall x.\, \varphi \ ,$$

where $x \in \mathcal{X}$. For readability, we lift the quantification over variables to the quantification over sets of variables and denote a maximal consecutive block of quantifiers of the same type $\forall x_1.\, \forall x_2. \cdots \forall x_n.\, \varphi$ by $\forall X.\, \varphi$ and $\exists x_1.\, \exists x_2. \cdots \exists x_n.\, \varphi$ by $\exists X.\, \varphi$, accordingly, where $X = \{x_1, \ldots, x_n\}$.

Given a subset of variables $X \subseteq \mathcal{X}$, an *assignment* of X is a function $\alpha : X \to \mathbb{B}$ that maps each variable $x \in X$ to either true (1) or false (0). When the domain of α is not clear from context, we write α_X. A partial assignment $\beta : X \to \mathbb{B} \cup \{\bot\}$ may additional set variables $x \in X$ to an undefined value $\bot$. We say that β is *compatible* with α, $\beta \sqsubseteq \alpha$ for short, if they have the same domains ($\mathrm{dom}(\alpha) = \mathrm{dom}(\beta)$) and $\alpha(x) = \beta(x)$ for all $x \in X$ where $\beta(x) \neq \bot$. For two assignments α and α' with disjoint domains $X = \mathrm{dom}(\alpha)$ and $X' = \mathrm{dom}(\alpha')$ we define the combination $\alpha \sqcup \alpha' : X \cup X' \to \mathbb{B}$ as $\alpha \sqcup \alpha'(x) = \begin{cases} \alpha(x) & \text{if } x \in X, \\ \alpha'(x) & \text{otherwise.} \end{cases}$ We define the *complement* $\overline{\alpha}$ to be $\overline{\alpha}(x) = \neg\alpha(x)$ for all $x \in \mathrm{dom}(\alpha)$. The *set of assignments* of X is denoted by $\mathcal{A}(X)$.

A quantifier $Q\, x.\, \varphi$ for $Q \in \{\exists, \forall\}$ *binds* the variable x in the *scope* φ. Variables that are not bound by a quantifier are called *free*. The set of free variables of formula φ is defined as $\mathit{free}(\varphi)$. The semantics of the satisfaction relation $\alpha \vDash \varphi$ is given as

$\alpha \vDash x$	if $\alpha(x) = 1$,
$\alpha \vDash \neg\varphi$	if $\alpha \nvDash \varphi$,
$\alpha \vDash \varphi \vee \psi$	if $\alpha \vDash \varphi$ or $\alpha \vDash \psi$,
$\alpha \vDash \varphi \wedge \psi$	if $\alpha \vDash \varphi$ and $\alpha \vDash \psi$,
$\alpha \vDash \exists y.\, \varphi$	if some $\alpha' : \{y\} \to \mathbb{B}$ satisfies $\alpha \sqcup \alpha' \vDash \varphi$, and
$\alpha \vDash \forall y.\, \varphi$	if all $\alpha' : \{y\} \to \mathbb{B}$ satisfy $\alpha \sqcup \alpha' \vDash \varphi$.

QBF satisfiability is the problem to determine, for a given QBF φ, the existence of an assignment α for the free variables of φ, such that the relation $\vDash$ holds.

The dependency set of an existentially (universally) quantified variable y, denoted by $\mathit{dep}(y)$, is the set X of universally (existentially) quantified variables x such that $Q\, y.\, \varphi$ is in the scope of x. A universal variable additionally depends on the free variables. A Boolean function $f : \mathcal{A}(X) \to \mathbb{B}$ maps *assignments* of X to

true or false. An assignment α over variables X can be identified by the conjunctive formula $\bigwedge_{x \in X | \alpha(x)=1} x \wedge \bigwedge_{x \in X | \alpha(x)=0} \neg x$. Similarly, Boolean functions can be identified by propositional formulas over the variables in their domain. Let $\varphi[f_{y_1}, \ldots, f_{y_n}]$ be the propositional formula where all quantifiers are removed and occurrences of variables $y_1, \ldots, y_n$ are replaced by their corresponding functions. A witness for a satisfiable QBF is a *Skolem function* $f_y : \mathcal{A}(dep(y)) \to \mathbb{B}$ for every existential variable y, such that for every assignment α of the universal variables $\alpha \wedge \varphi[f_{y_1}, \ldots, f_{y_n}]$ is satisfiable. For unsatisfiable QBFs, the witnesses are defined dually and called *Herbrand functions*.

A *closed* QBF is a formula without free variables. Closed QBFs are either true or false. A formula is in prenex form, if the formula consists of a quantifier prefix followed by a propositional formula. Every QBF can be transformed into a closed QBF and into prenex form while maintaining satisfiability.

A *literal* l is either a variable $x \in X$, or its negation $\neg x$. Given a set of literals $\{l_1, \ldots, l_n\}$, the disjunctive combination $(l_1 \vee \ldots \vee l_n)$ is called a *clause* and the conjunctive combination $(l_1 \wedge \ldots \wedge l_n)$ is called a *cube*. Given a literal l, the polarity of l, $sign(l)$ for short, is 1 if l is positive and 0 otherwise. The variable corresponding to l is defined as $var(l) = x$ where $x = l$ if $sign(l) = 1$ and $x = \neg l$ otherwise.

A QBF is in negation normal form (NNF) if negation is only applied to variables. Every QBF can be transformed into NNF by at most doubling the size of the formula and without introducing new variables. For formulas in NNF, we treat literals as atoms.

3 Solving Quantified Boolean Formulas

For this section, we assume an arbitrary quantified formula $Q\,X_1 \ldots Q\,X_n . \varphi$ with quantifier prefix $Q\,X_1 \ldots Q\,X_n$ and propositional body φ in NNF.

We use a generic solving function $\textsc{sat}(\theta, \alpha)$ for propositional formula θ under assumptions α, that returns whether $\theta \wedge \alpha$ is satisfiable and either a satisfying assignment α_V for variables $V \subseteq free(\theta)$ or a partial assignment $\beta_{failed} \sqsubseteq \alpha$ such that $\theta \wedge \beta_{failed}$ is unsatisfiable. Further, we define SAT_Q to be SAT if $Q = \exists$ and UNSAT otherwise (UNSAT_Q analogously).

3.1 Assignment Based Solving

Algorithm 1 shows a recursive solving algorithm that communicates variable assignments through the quantifier levels of the QBF.

Every quantifier $\exists X . \psi$ has a propositional formulas θ_X that is initialized with the body φ of the quantified formula ($\theta_X = \neg\varphi$ for universal quantifiers). In every iteration of the loop, an X assignment is generated in line 3. In the innermost quantifier, we optimize the assignment α' by generating a partial assignment β'_{failed} of α' that is sufficient to falsify $\neg\theta_X$ [33]. Consequently, this partial assignment β'_{failed} is sufficient to satisfy θ_X. If ψ is quantified, we recursively solve ψ under the assignment α'. In the positive case ($result = \text{SAT}_Q$), we

Algorithm 1 Assignment Based Algorithm

```
1:  procedure ASSIGNMENT-QBF(Q X. ψ, α)
2:      while true do
3:          result, α_X, β_failed ← SAT(θ_X, α)                    ▷ β_failed ⊑ α
4:          if result = UNSAT then
5:              return UNSAT_Q, β_failed
6:          α' ← α_X ⊔ α
7:          if ψ is propositional then
8:              result, _, β'_failed ← SAT(¬θ_X, α')               ▷ result = UNSAT, β'_failed ⊑ α'
9:              return SAT_Q, β'_failed
10:         result, β ← ASSIGNMENT-QBF(ψ, α')                      ▷ β ⊑ α'
11:         if result = SAT_Q then
12:             return SAT_Q, β
13:         θ_X ← θ_X ∧ ¬β
```

return the obtained witness β while in the negative case, we refine a clause that excludes the counterexample assignment before reentering the loop. If there is no satisfying assignment in line 3, the algorithm returns a partial assignment $\beta_{failed} \sqsubseteq \alpha$ that witnesses the unsatisfiability of the query.

A simple inductive argument shows that, given a QBF, the algorithm terminates with the correct result.

3.2 Abstraction Based Solving

Refinements based on variable assignments are in many cases not optimal. This insight led to algorithms that generalizes from assignment based refinement to clause based refinement [28,37] for QBFs given in conjunctive normal form. The idea in these works is that multiple assignments may lead to the satisfaction of the same clauses, hence, instead of communicating assignments, the information whether a clause is satisfied or not is communicated between quantifier levels. While in the assignment based algorithm, a refinement (line 13) states that at least one assignment of β must be different in the next iteration, the refinement in the clause based algorithm states that from a set of clauses/cubes one element must be satisfied (depending on the quantifier).

In the following, we propose another generalization to negation normal form formulas, i.e., we base the communication on the satisfaction of individual subformulas. To facilitate working with arbitrary Boolean formulas, we start with introducing additional notation. Let $\mathcal{B}$ be the set of Boolean formulas and let $sf(\psi) \subset \mathcal{B}$ ($dsf(\psi) \subset \mathcal{B}$) be the set of (direct) subformulas of ψ (note that $\psi \in sf(\psi)$ but $\psi \notin dsf(\psi)$). For a propositional formula ψ, $type(\psi) \in \{lit, \vee, \wedge\}$ returns the Boolean connector if ψ is not a literal. For example, given $\psi = (x_1 \vee (\overline{x_1} \wedge x_2))$, $sf(\psi) = \{(x_1 \vee (\overline{x_1} \wedge x_2)), x_1, (\overline{x_1} \wedge x_2), \overline{x_1}, x_2\}$, $dsf(\psi) = \{x_1, (\overline{x_1} \wedge x_2)\}$, and $type(\psi) = \vee$.

To communicate the satisfaction of subformulas, we introduce two special types of literals which we call *interface literals*. The valuation of those liter-

als are communicated between the quantifier levels. For a given quantifier $Q\,X$ and a subformula ψ, the T literal t_ψ represents the assignments made by the outer quantifiers while the B literal b_ψ represents the assignments from the current quantifier including assumptions on the satisfaction of subformulas by inner quantifiers. Thus, t_ψ is true if ψ is satisfied by the outer quantifiers and a valuation that sets b_ψ to true indicates that ψ is satisfied by the quantifier $Q\,X$. Before going in more detail on the abstraction, we introduce the basic algorithm first.

Recursive CEGAR Algorithm. Algorithm 2 shows the abstraction based algorithm ABSTRACTION-QBF. On a high level, the algorithm has the same structure as ASSIGNMENT-QBF as it does recursion on the quantifier blocks. Unlike to previous approaches that utilize SAT solvers [26, 28, 37, 40], we keep for every scope a *dual abstraction* that is used for optimization of abstraction entries and translating interface literals.

Given a quantifier $Q\,X$, the sets T_X and B_X contain the T and B literals corresponding to this quantifier. When translating a B literal to a T literal, we use the same index, e.g., in line 9, the B literals $b_{\psi'}$ are translated to T literals $t_{\psi'}$ of the inner scope. We initialize θ_X with the abstraction described below, which is a propositional formula over variables in X, T and B literals. The innermost quantifier does not contain B literals. The *dual abstraction* $\overline{\theta}_X$ is defined as the abstraction for $\overline{Q}\,X$.

In every iteration of the while loop, a B literal assignment α_{B_X} is generated according to the outer T assignment α_{T_X} (line 3). Before the recursive call, α_{B_X} is translated to a T_Y assignment α_{T_Y} for the inner quantifier $\overline{Q}\,Y$. If the recursive call is UNSAT (w.r.t. the current quantifier), the abstraction θ_X is refined by excluding the counterexample β_{T_Y}, i.e., in the following iterations one of the $b_{\psi'}$ such that $\beta_{T_Y}(t_{\psi'}) = 1$ must be set to true. Due to the negation during translation, those $b_{\psi'}$ were set to false in α_{B_X}. In case the recursive call is SAT, we update the dual abstraction accordingly and use it to optimize and translate the witness β_{T_Y} from a T_Y assignment to a T_X assignment. If the query in line 3 fails, UNSAT_Q together with an assignment $\beta_{\mathit{failed}} \sqsubseteq \alpha_{T_X}$ witnessing the unsatisfiability is returned.

To ensure correctness, we need requirements on the abstraction being used. Given a quantifier $\exists X$, we say that a subformula ψ is good if it is not yet falsified ($\mathit{type}(\psi) = \wedge$), respectively satisfied ($\mathit{type}(\psi) = \vee$). For every quantifier $Q\,X$ and B literal b_ψ it must hold that if b_ψ is set to true then ψ is good for quantifier $Q\,X$, which gives us proper refinement semantics (line 15). The same property holds for T literals t_ψ. Further, for a quantifier alternation $Q\,X.\,\overline{Q}\,Y$, it must hold that the set of outer B literals B_X matches the set of inner T literals T_Y to enable the translation in line 9. Combining these properties gives us that a good subformula of quantifier $Q\,X$ is a bad subformula of quantifier $\overline{Q}\,Y$ and vice versa. Termination then follows from progress due to refinements (line 15) and correctness can be showed by induction over the quantifier blocks.

Algorithm 2 Abstraction Based Algorithm

1: **procedure** ABSTRACTION-QBF($Q\,X.\,\psi, \alpha_{T_X}$)
2: **while** true **do**
3: $\mathit{result}, \alpha_X \sqcup \alpha_{B_X}, \beta_{\mathit{failed}} \leftarrow \text{SAT}(\theta_X, \alpha_{T_X})$ ▷ $\beta_{\mathit{failed}} \sqsubseteq \alpha_{T_X}$
4: **if** $\mathit{result} = \text{UNSAT}$ **then**
5: **return** $\text{UNSAT}_Q, \beta_{\mathit{failed}}$
6: **else if** ψ is propositional **then**
7: $\mathit{result}, _, \beta_{\mathit{failed}} \leftarrow \text{SAT}(\overline{\theta}_X, \alpha_X \sqcup \overline{\alpha_{T_X}})$ ▷ $\mathit{result} = \text{UNSAT}$
8: **return** $\text{SAT}_Q, \overline{\beta_{\mathit{failed}}}$ ▷ $\overline{\beta_{\mathit{failed}}} \sqsubseteq \alpha_{T_X}$
9: Let $\psi = \overline{Q}\,Y.\,\psi^*$ and define α_{T_Y} s.t. $\alpha_{T_Y}(t_{\psi'}) = \neg\alpha_{B_X}(b_{\psi'})$ for all $t_{\psi'} \in T_Y$
10: $\mathit{result}, \beta_{T_Y} \leftarrow \text{ABSTRACTION-QBF}(\psi, \alpha_{T_Y})$ ▷ $\beta_{T_Y} \sqsubseteq \alpha_{T_Y}$
11: **if** $\mathit{result} = \text{SAT}_Q$ **then**
12: $\overline{\theta}_X \leftarrow \overline{\theta}_X \wedge (\bigvee_{b_{\psi'} \in B_X \mid \beta_{T_Y}(t_{\psi'})=1} b_{\psi'})$
13: $\mathit{result}, _, \beta_{\mathit{failed}} \leftarrow \text{SAT}(\overline{\theta}_X, \alpha_X \sqcup \overline{\alpha_{T_X}})$ ▷ $\mathit{result} = \text{UNSAT}$
14: **return** $\text{SAT}_Q, \overline{\beta_{\mathit{failed}}}$ ▷ $\overline{\beta_{\mathit{failed}}} \sqsubseteq \alpha_{T_X}$
15: $\theta_X \leftarrow \theta_X \wedge (\bigvee_{b_{\psi'} \in B_X \mid \beta_{T_Y}(t_{\psi'})=1} b_{\psi'})$

Abstraction. Before giving the formal definition, we give an intuition for the abstraction and its interference with the algorithm.

Example 1. Consider the example $\Phi_{ex} = \forall x.\, \exists y.\, \overbrace{(x \vee \underbrace{(\overline{x} \wedge y)}_{\psi_2})}^{\psi_1}$ and its negation $\neg\Phi_{ex} = \exists x.\, \forall y.\, (\overline{x} \wedge (x \vee \overline{y}))$. The abstraction for the universal quantifier $\forall x$ is $\theta_{\{x\}} = (b_{\psi_1})(\neg b_{\psi_1} \vee \overline{x})(\neg b_{\psi_2} \vee x)$ and the abstraction for the existential quantifier $\exists y$ is $\theta_{\{y\}} = (t_{\psi_1} \vee b_{\psi_2})(\neg b_{\psi_2} \vee t_{\psi_2})(\neg b_{\psi_2} \vee y)$. The only satisfying assignment $\alpha_{B_{\{x\}}}$ for $\theta_{\{x\}}$ is $\alpha_{B_{\{x\}}}(b_{\psi_1}) = 1$ and $\alpha_{B_{\{x\}}}(b_{\psi_2}) = 0$, i.e., the universal quantifier can satisfy ψ_1 in $\neg\Phi_{ex}$, while assuming that ψ_2 will be satisfied by the inner level (as b_{ψ_2} is set to false). This assignment $\alpha_{B_{\{x\}}}$ is translated into T assignment $\alpha_{T_{\{y\}}}$ of the T literals in $\varphi_{\{y\}}$ by negation ($\alpha_{T_{\{y\}}}(t_{\psi_1}) = 0$ and $\alpha_{T_{\{y\}}}(t_{\psi_2}) = 1$). At the existential quantifier, the abstraction $\theta_{\{y\}}$ is solved under the assumption $\alpha_{T_{\{y\}}}$, resulting in a satisfiable query. The assumption $\alpha_{T_{\{y\}}}(t_{\psi_2}) = 1$ is needed to satisfy $\theta_{\{y\}}$ (assumptions $\neg t_{\psi_2}$ and y are enough to make the dual abstraction $\overline{\theta}_{\{y\}} = (b_{\psi_1})(\neg b_{\psi_1} \vee t_{\psi_1})(\neg b_{\psi_1} \vee b_{\psi_2})(\neg b_{\psi_2} \vee t_{\psi_2} \vee \overline{y})$ unsatisfiable, line 7), thus, the partial assignment β with $\beta(t_{\psi_2}) = 1$ is returned (line 8). The following refinement forces that b_{ψ_2} must be set to true in the next iteration, i.e., $\theta_{\{x\}} = \theta_{\{x\}} \wedge b_{\psi_2}$. This depletes all possible assignments of the universal quantifier, thus, proving that the instance is true.

We now give a formal definition of the abstraction. Given a propositional formula φ in NNF and a scope $\exists X$, we build the following propositional formula in conjunctive normal form representing the abstraction $\theta_X = \mathit{out}(\varphi) \wedge \bigwedge_{\psi \in sf(\varphi) \wedge type(\psi) \neq lit} \mathit{enc}(\psi)$ for this scope, where *out* encodes that φ must hold and *enc* defines a CNF formula that encodes the truth of subformula ψ with respect to the valuations of the current, inner and outer scopes represented by

B and T literals, respectively:

$$enc(\psi) = \begin{cases} \bigwedge\limits_{\psi' \in dsf(\psi)} (\neg b_\psi \vee enc_\psi(\psi')) & \text{if } type(\psi) = \wedge \\ \neg b_\psi \vee \bigvee\limits_{\psi' \in dsf(\psi)} enc_\psi(\psi') & \text{if } type(\psi) = \vee \end{cases}$$

$$enc_\psi(\psi') = \begin{cases} \psi' & \text{if } type(\psi') = lit \wedge var(\psi') \in X \\ t_\psi & \text{if } type(\psi') = lit \wedge var(\psi') \text{ bound by outer scope} \\ \neg b_\psi & \text{if } type(\psi') = lit \wedge var(\psi') \text{ bound by inner scope} \\ b_{\psi'} & \text{if } type(\psi') \neq lit \wedge \psi' \text{ only influenced by current or outer scope} \\ \bot & \text{otherwise} \end{cases}$$

$$enc_\vee(\psi) = \bigvee_{\substack{\psi' \in dsf(\psi) \\ type(\psi') = lit}} enc_\psi(\psi') \quad \vee \bigvee_{\substack{\psi' \in dsf(\psi) \\ type(\psi') = \wedge}} b_{\psi'} \quad \vee \bigvee_{\substack{\psi' \in dsf(\psi) \\ type(\psi') = \vee}} enc_\vee(\psi')$$

$$out(\psi) = \begin{cases} b_\psi & \text{if } type(\psi) = \wedge \\ enc_\vee(\psi) & \text{if } type(\psi) = \vee \end{cases}$$

An undefined result $\bot$ form $enc_\psi(\psi')$ means that the subformula ψ' is ignored in the encoding $enc(\psi)$. The abstraction of a scope $\forall X$ is defined as the existential abstraction for $\neg\varphi$. The dual abstraction $\overline{\theta}_X$ is defined as the abstraction for $\neg Q\, X$. Note that not every B literal that is used in the abstraction may be exposed as an interface literal. For this abstraction we define

$$B_X = \{b_\psi \mid \psi \in sf(\varphi) \wedge type(\psi) \neq lit \wedge \psi \text{ contains variable of inner scope}\},$$
$$T_X = \{t_\psi \mid \psi \in sf(\varphi) \wedge type(\psi) \neq lit \wedge \psi \text{ contains variable of outer scope}\}.$$

These sets can be further reduced as described in the next section.

Optimizations. There are a number of optimizations that can be applied to the basic abstraction algorithm. First, the SAT solver used in line 3 does not know about the semantics of the T and B literals, hence, the returned assignment may not be optimal. Depending on the solver used, there is limited influence on the solving, like setting default polarities, control decision heuristics, etc. A SAT solver independent solution to this problem is a pass after the SAT call that adjusts the B literal assignment α_{B_X} according to the chosen variable assignment. This is possible since in NNF (as in CNF), the optimization objective for the existential and universal quantifiers is to maximize and minimize the number of satisfied subformulas, respectively. More sophisticated optimization could be achieved by taking the formula structure into account or even modify the variable assignment α_X to optimize the B literal assignment.

The abstraction can be further improved to reduce the number of interface literals. By computing which subformula is influenced by which scopes, one can combine a whole formula subtree into a single B literal if all subformulas in the subtree are only influenced by the current or outer scopes. We depict such a situation in Fig. 1.

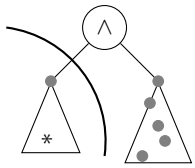

Fig. 1. Visualization of the scope influence optimization: In this picture, the subtree $(*)$ is only influenced by outer scope S, so we collapse the whole subtree into a single T literal (indicated by grey dots) in the inner scopes. For the abstraction at scope S, we correspondingly introduce a just single B literal for the output of the subtree.

Preprocessing. Due to the richer structure of formulas compared to CNF, it is more difficult to implement preprocessing techniques and we consider only basic ones. We normalize the formula by combining consecutive subformulas with the same operator, removing tautologies and forced literals. As in NNF negation appears only in front of variables, detection of pure literals [15] is still efficient. Since usually circuits are used to represent formulas, sophisticated circuit minimization techniques would likely benefit solving although we do not make use of them.

4 Certification

Certification is an essential complement of QBF solving. Certification amounts to extracting witnessing functions from a QBF, either *Skolem* functions for true QBFs or *Herbrand* functions for false QBFs. Not only does it allow to verify the solver result, but also can the resulting functions be used in the context of the application. More recently, QBF proofs were used to establish differentiations between different kind of PCNF solving methods [4, 27]. We use some of these results to separate our algorithm from other QBF solving algorithms in Sec. 5. The main result of this section is a proof format for our abstraction algorithm and an efficient algorithm to transform proof traces into Boolean functions.

Proof Format. To extract a witness from a run of ABSTRACTION-QBF, we need to remember *situations* and *reactions*, represented by assignments to T and X, that were satisfiable for the respective quantifier. Hence, the proof $\mathcal{P}$ consists of a sequence of pairs $\langle \boldsymbol{t}, \boldsymbol{x} \rangle \in (2^T \times 2^X)^*$ and these pairs can be obtained from the algorithm by the result β_{failed} of the query to the dual abstraction $\overline{\theta}_X$ in lines 7 and 13. As an immediate consequence, the number of pairs in the proof trace is linear in the number of iterations of the algorithm. Further, the proof format is *independent* of the abstraction that is used for the algorithm.

As the function extraction depends on the abstraction, we define a function $\mathcal{C} : T \times X \to \mathcal{B}(X')$ which maps a T literal and a scope X to a Boolean formula over the variables X' of the outer scopes (with respect to X). Intuitively, $\mathcal{C}(t, X)$ describes these assignments that enable T in the abstraction of scope X.

Example 2. Consider again our earlier example formula $\Phi_{ex} = \forall x.\, \exists y.\, \overbrace{(x \vee \underbrace{(\overline{x} \wedge y)}_{\psi_2})}^{\psi_1}$.

It holds that $\mathcal{C}(t_{\psi_1}, \{x\}) = x$ and $\mathcal{C}(t_{\psi_2}, \{x\}) = \overline{x}$.

Function Extraction. Prior to the function extraction, we filter out those pairs from the proof $\mathcal{P}$ that correspond to the variables that are dependencies and do not describe a function, i.e., universal variables for true QBFs and existential variables for false QBFs. The remaining proof consists of pairs $\langle \boldsymbol{t}, \boldsymbol{x} \rangle$ where $\boldsymbol{x}$ is a set of variables that were assigned positively and $\bigwedge_{t \in \boldsymbol{t_i}} \mathcal{C}(t_i, X)$ is a formula that represents the situation were the response $\boldsymbol{x}$ is correct. Let $\langle \boldsymbol{t_1}, \boldsymbol{x_1} \rangle \ldots \langle \boldsymbol{t_n}, \boldsymbol{x_n} \rangle$ be the pairs corresponding to scope X and let $x \in X$ be some variable, the function $f_x : \mathcal{A}(dep(x)) \to \mathbb{B}$ is defined as

$$f_x(dep(x)) \equiv \bigvee_{i=1}^{n} \left((x \in \boldsymbol{x_i}) \wedge \bigwedge_{t \in \boldsymbol{t_i}} \mathcal{C}(t_i, X) \wedge \bigwedge_{j<i} \neg \bigwedge_{t \in \boldsymbol{t_j}} \mathcal{C}(t_j, X) \right). \tag{1}$$

This construction is similar to previous extraction algorithms, including [4, 37]. The definition of $\mathcal{C}$ allows that f_x may depend on variables in outer scopes corresponding to functions instead of dependencies. By replacing those variables with their extracted functions, one can make sure that f_x depends only on $dep(x)$. The size of f_x, measured in terms of distinct subformulas, is linear in the number of pairs and, hence, linear in the size of the proof.

Example 3. The proof trace for example Φ_{ex} is $\langle \{t_{\psi_2}\}, \{y\} \rangle$ and the Skolem function is $f_y(x) \equiv \mathcal{C}(t_{\psi_2}, \{x\}) \equiv \overline{x}$.

Theorem 1. *Given a QBF Φ and proof trace $\mathcal{P}$. The run time of the function extraction algorithm is in $\mathcal{O}(|\mathcal{P}|)$. The size of the resulting functions is linear in the size of $\mathcal{P}$.*

Certification. A *certificate* is a representation of all functions that witnesses the result of the QBF. A certificate is correct, if two conditions are satisfied: it is (1) functionally correct and (2) well-formed. Functionally correctness can be checked by a propositional SAT query to $\neg\varphi$ (respectively φ for false QBFs) where every occurrence of a function variable x is replaced by the function f_x. The unsatisfiability of this query witnesses functionally correctness. The well-formedness criterion concerns the representation of the certificate, usually as a circuit, and requires that the representation of a function depends only on its dependencies. One can further differentiate *syntactical* and *semantical* well-formedness. A certificate is syntactical ill-formed if a non-dependency is reachable from the output of a function. A certificate is semantical ill-formed if a valuation change of a set of non-dependencies changes the valuation of a function. Our function extraction guarantees syntactical well-formedness and therefore any further circuit simplification guarantees at least semantical well-formedness.

5 Comparison to other Solving Algorithms

In this section we compare ABSTRACTION-QBF to well-established solving methods. For comparing the runtime of ABSTRACTION-QBF and ASSIGNMENT-QBF, we use the number of iterations of the while loops as a measurement.

First, we show that there is an exponential gap between ABSTRACTION-QBF and ASSIGNMENT-QBF. We use the family QPARITY_n of QBFs that was used to show separation between Q-resolution and expansion based solving [5]. QPARITY_n is defined as $\exists X. \forall z. \underbrace{(z \vee \text{PARITY}(X))}_{\psi_1} \wedge \underbrace{(\neg z \vee \neg \text{PARITY}(X))}_{\psi_2}$ with $|X| = n$ for every $n > 1$ and $\text{PARITY}(X) = \bigoplus_{x \in X} x$. QPARITY_n is false for every bound and the unique Herbrand function f_z is defined as $f_z = \text{PARITY}(X)$.

Lemma 1. *Every run of* ASSIGNMENT-QBF *on* QPARITY_n *needs exponentially many iterations in* n*.*

Proof. We show that every refinement at the existential quantifier excludes exactly one function entry of f_z. As there are 2^n function entries, ASSIGNMENT-QBF needs an exponential number of iterations.

Let α_X be an arbitrary assignment. Assume that $\text{PARITY}(X)$ is true under this assignment (the other case is analogous). At the universal quantifier, the query $(\neg z \wedge \neg \text{PARITY}(X)) \vee (z \wedge \text{PARITY}(X))$ is satisfied under assumption α_X by choosing $z = 1$. As flipping the value of already one $x \in X$ would make $\text{PARITY}(X)$ false, one need the complete assignment $\alpha_X \sqcup \alpha_z$ to falsify $(z \vee \text{PARITY}(X)) \wedge (\neg z \vee \neg \text{PARITY}(X))$ (line 8 of Algorithm 1). Hence, the assignment α_X is returned by the universal level and the refinement eliminates exactly one function entry. □

Lemma 2. *Every run of* ABSTRACTION-QBF *on* QPARITY_n *needs constant number of iterations.*

Proof. The abstraction $\theta_X \equiv (\neg b_{\psi_1} \vee enc_{\text{PARITY}(X)})(\neg b_{\psi_2} \vee enc_{\neg \text{PARITY}(X)})$ has only two relevant B literals, namely b_{ψ_1} and b_{ψ_2} and the existential quantifier has to set b_{ψ_1} or b_{ψ_2} to false (it may be the case that both are set to false). The universal quantifier needs just one T literal set to true to satisfy $\theta_{\{z\}} \equiv (b_{\psi_1} \vee b_{\psi_2})(\neg b_{\psi_1} \vee \neg z)(\neg b_{\psi_1} \vee t_{\psi_1})(\neg b_{\psi_2} \vee z)(\neg b_{\psi_2} \vee t_{\psi_2})$, hence, either b_{ψ_1} or b_{ψ_2} is refined (line 15 of Algorithm 2). The abstraction θ_X becomes unsatisfiable after the next iteration, which is analogous to the first one. □

We use a simulation argument on the proof traces to show that ABSTRACTION-QBF is at least as good as ASSIGNMENT-QBF.

Lemma 3. *Given an arbitrary QBF* Φ*. Every proof trace of* ASSIGNMENT-QBF *on* Φ *can be polynomially simulated by a proof trace of* ABSTRACTION-QBF*.*

Proof (Sketch). A proof trace $\mathcal{P}$ for ASSIGNMENT-QBF can be defined by a sequence of pairs $\langle \boldsymbol{x'}, \boldsymbol{x} \rangle$ where for a given scope X, $\boldsymbol{x'}$ is a set (cube) of variables from outer scopes that indicates $\boldsymbol{x}$ is a satisfying assignment under this condition.

Given such a trace $\mathcal{P}$, we translate it into a proof trace of ABSTRACTION-QBF by computing for each pair $\langle \boldsymbol{x'}, \boldsymbol{x} \rangle$ which interface T literals $\boldsymbol{t}$ are satisfied by $\boldsymbol{x'}$, resulting in the pair $\langle \boldsymbol{t}, \boldsymbol{x} \rangle$. By construction, the resulting proof trace $\mathcal{P}'$ is correct and $|\mathcal{P}| = |\mathcal{P}'|$.

Theorem 2.

1. ASSIGNMENT-QBF *does not simulate* ABSTRACTION-QBF.
2. ABSTRACTION-QBF *simulates* ASSIGNMENT-QBF.

We compare ABSTRACTION-QBF to expansion based algorithms in the following. To do so, we use a family of formulas

$$\exists e_1 \forall u_1 \exists c_1 c_2 \ldots \exists e_n \forall u_n \exists c_{2n-1} c_{2n}.$$
$$\bigwedge_{i \in \{1,\ldots,n\}} \Big(\big((e_i \vee u_i) \to c_{2i-1} \big) \big((\overline{e_i} \vee \overline{u_i}) \to c_{2i} \big) \Big) \wedge \bigvee_{i \in \{1,\ldots,2n\}} \overline{c_i} \tag{2}$$

that was used to show that ∀Exp+Res does not simulate Q-resolution [27]. The refutation of (2) needs to expand each variable in both polarities at all depths [27], hence, there is no polynomial refutation in expansion based solvers like RAReQS [26].

Lemma 4. *There exists a run of* ABSTRACTION-QBF *on formula (2) that has a linear number of iterations with respect to n.*

Proof (Sketch). We consider the NNF translation (3) of (2). W.l.o.g. we can consider the quantifier prefix as n blocks of quantifiers $B_1, \ldots, B_n$ as the subformula D_i contains only variables from block B_i for every $i \in \{1, \ldots, n\}$, i.e., the only subformula shared between quantifiers is C.

$$\overbrace{\exists e_1 \forall u_1 \exists c_1 c_2}^{B_1} \ldots \overbrace{\exists e_n \forall u_n \exists c_{2n-1} c_{2n}}^{B_n}.$$
$$\bigwedge_{i \in \{1,\ldots,n\}} \underbrace{\Big(\big(\underbrace{(\overline{e_i} \wedge \overline{u_i})}_{E_{i,1}} \vee c_{2i-1} \big) \big(\underbrace{(e_i \wedge u_i)}_{E_{i,2}} \vee c_{2i} \big) \Big)}_{D_n} \wedge \underbrace{\bigvee_{i \in \{1,\ldots,2n\}} \overline{c_i}}_{C} \tag{3}$$

For every quantifier $\exists e_i$, either cube $E_{i,1}$ or $E_{i,2}$ must be falsified. Hence, by using $u_i = \neg e_i$, the inner quantifier $\forall u_i$ can falsify both cubes $E_{i,1}$ and $E_{i,2}$, forcing c_{2i-1} and c_{2i} to be true. Further, and unlike to the expansion based setting, the situations where $\{e_i \to 1, u_1 \to 0\}$ and $\{e_i \to 0, u_1 \to 1\}$ are the same with respect to our abstraction (both cubes are falsified), hence, these situations fall together.

Using this, we can construct a run that proceeds in this way until the innermost quantifier $\exists c_{2n-1}, c_{2n}$. The corresponding query is UNSAT as clause C is not satisfied yet, i.e., requires $(\overline{c_{2n-1}} \vee \overline{c_{2n}})$, and c_{2n-1} as well as c_{2n} are forced to be true. The failed assumptions of this query correspond to the T literals of $E_{n,1}$,

$E_{n,2}$, and C. Due to the optimization in the dual abstraction of quantifier $\forall u_n$, this reduces to either $E_{n,1}$, C or $E_{n,2}$, C if $e_n = 1$ or $e_n = 0$, respectively. The refinement on $\exists e_n$ forces e_n to satisfy the previous falsified cube, i.e., the value of e_i is swapped and the algorithm proceeds as above, but this time returns the other refinement $E_{n,2}$, C or $E_{n,1}$, C. This establishes that C must be satisfied at this level, thus, eliminating one block B_n.

6 Experimental Results

We implemented the algorithms described in this paper and their optimizations in a solver called QuAbS[1] (Quantified Abstraction Solver) that takes QBFs in the standard format QCIR [36]. As the underlying SAT solver, we use PicoSAT [6] and MiniSat [14]. We compare QuAbS against the two publicly available QBF solvers that support the QCIR format, namely RAReQS [26] and GhostQ [30]. For our experiments, we used a machine with a 3.6 GHz quad-core Intel Xeon processor and 32 GB of memory. The timeout and memout were set to 10 minutes and 8 GB, respectively.

Benchmarks. We evaluate the solver on three novel benchmark sets that are designed to take advantage of the structural problem definition. The *syntcomp* benchmark set was created from the safety benchmarks of SYNTCOMP 2014 [25]. In this terminology, X and X' are the state and next state bits, I is the environment input, O is the system output, $T \subseteq \mathcal{B}(X, I, O, X')$ is the transition relation, and $W \subseteq \mathcal{B}(X)$ denotes the winning region. Starting from all safe states $W = \varphi_{\text{safe}}$, we iteratively solve the query $\forall X, I. \exists O, X'. W \rightarrow (T \wedge W')$ and adapting $W = W \wedge \neg \alpha_X$ for false queries [10], i.e., we exclude states represented by α_X form where the environment can force the system into a bad state. If the query becomes satisfiable, the benchmark is realizable, if W excludes the initial state φ_{init}, the benchmark is unrealizable. The *bounded-synthesis* benchmark set was generated from the synthesis benchmarks bundled with Acacia+ [11] using a QBF encoding of the bounded synthesis approach [18]. The *petri-games* benchmark set was generated using a variant of the bounded synthesis approach for Petri games [16] on the benchmark set introduced by Finkbeiner et al. [17]. All those benchmarks have in common that it is possible to directly build implementations from satisfiable queries. The *qbfgallery2014* and *eval2012* are originally CNF benchmark sets that were translated to QCIR using a conversion tool provided by the QBFGallery organizers.

Evaluation. Table 1 shows the result of the evaluation broken down by benchmark families. Overall, QuAbS is very competitive. It clearly outperforms the assignment based approach, confirming the theoretical results established in the previous section. Compared to the established solver, the results are mixed. There are benchmarks where QuAbS outperforms both RAReQS and GhostQ like the Petri games family as well as benchmarks where RAReQS (bounded

[1] Available at `https://www.react.uni-saarland.de/tools/quabs/`

Table 1. Number of instances of the benchmark sets solved within 10 minutes.

Family	total	QuAbS		assignment		RAReQS	GhostQ	
		PicoSAT	MiniSat	PicoSAT	MiniSat		plain	CEGAR
syntcomp	160	154	**156**	135	138	109	**156**	147
bounded-synt.	330	281	276	261	237	**322**	278	319
petri-games	244	**213**	206	161	137	174	125	161
structural	734	**648**	638	557	512	605	559	627
qbfgallery2014	699	290	241	202	122	241	316	**333**
eval2012	337	177	166	98	77	101	161	**187**
translated	1036	467	407	300	199	342	477	**520**
total	1770	1115	1045	857	711	947	1036	**1147**

Table 2. Number of solved instances with certification enabled, the number of verified instances, average certificate sizes, and a breakdown of the relative times needed to solve and verify certificates by benchmark families.

Family	total	variant	#solved	#verified	avg. size	solving	verification
syntcomp	160	PicoSAT	154	142	39,900	19.7%	80.3%
		MiniSat	156	138	15,264	8.9%	91.1%
bounded-synthesis	330	PicoSAT	281	278	22,209	83.1%	16.9%
		MiniSat	276	270	54,865	45.3%	54.7%
petri-games	244	PicoSAT	213	204	129,542	42.7%	57.3%
		MiniSat	206	183	172,738	30.6%	69.4%
qbfgallery2014	699	PicoSAT	285	262	1,200,879	79.2%	20.3%
		MiniSat	238	208	2,742,076	64.0%	36.0%
eval2012	337	PicoSAT	177	161	523,123	48.7%	51.3%
		MiniSat	166	147	477,519	54.0%	46.0%

synthesis) and GhostQ (translated benchmarks) have an advantage. In the internal comparison, the variant using PicoSAT has a small lead over the variant using MiniSat, however, both variants solve some benchmarks that the other does not. Combining both leads to 659 and 505 solved structural and translated benchmarks, respectively.

Certification. We implemented the certification approach described in Sec. 4, but instead of generating proof traces, we build the certificates (repesented as AIGER files) within the solving loop. This allows us to build candidate Skolem and Herbrand functions in parallel during solving and minimizes the overhead due to certification. In the verification step, we build a SAT query that encodes the functional correctness property and solve this query with MiniSat. The size of a certificate is given in the number of AND gates in the AIGER file.

The overall effect of the certification approach on the running times is negligible (less than 1% increase for both variants) which confirms our goal that the combination of solving and certification can be implemented efficiently. Table 2

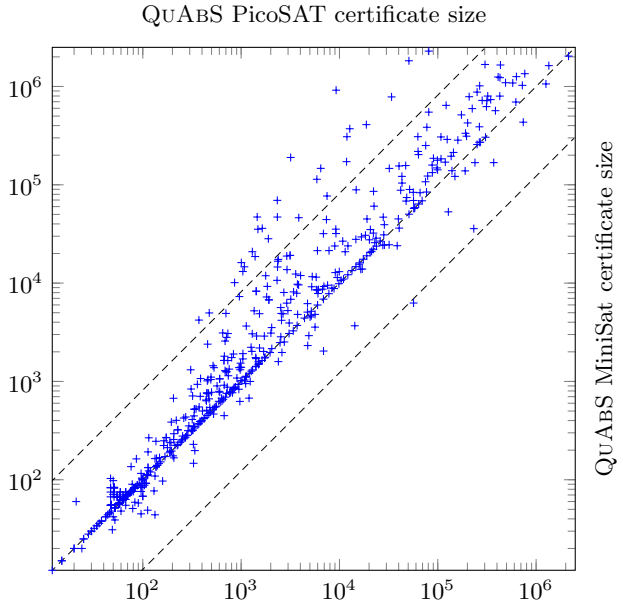

Fig. 2. The size of certificates computed by QuAbS depending on SAT solver on the structural benchmarks.

shows the results of the certification run. Compared to Table 1 only in the gbf-gallery2014 benchmark set there were some instances that could not be solved within timeout when certification was enabled. In most benchmarks, the time spent verifying the result exceeded the solving times. Consequently, not all of the solved instance could also be verified within timeout. We did not expect the verification to be a bottleneck, for future results we want to try more sophisticated SAT solvers reduce the time needed for verification.

The spike in the average size in the translated benchmarks is partially due to the (incomplete) translation, i.e., not every auxiliary CNF variable was detected and influences the certificate size negatively. To further reduce the size of certificates, we plan to employ available circuit minimization techniques.

Figure 2 compares the size produced by both solver variants on the commonly solved instances of the structural benchmarks. The observation correlates with Table 1 and 2 in that it shows a small advantage of the PicoSAT version, especially in seemingly more difficult instances, that are instances with large certificates.

7 Related Work

There are other solving techniques that use structural information, but they are conceptually very different, including DPLL like [15, 22, 30] and expansion [26, 31,35]. SAT solvers have been used to solve quantifier alternations [26,28,37,40],

also in the context of synthesis [10]. We extend work on QBF solving techniques that communicate the satisfaction of clauses through a recursive refinement algorithm [28, 37] that were limited to conjunctive normal form. Further, the maintenance of a dual abstraction for optimization is new in this context and the certification approach is different than [37]. Certification has been considered in the context of CNF solving techniques [2, 24, 34, 37] but we are not aware of another work considering certification in the more general setting. The duality of circuit based QBF solving has been used to enhance search based CNF solvers [21, 23] but our use of a dual abstraction during solving is inspired by recent SMT solving enhancements [33].

8 Conclusion

We presented a new QBF solving algorithm that exploits the structure in the propositional formula. Further, we defined a certification format suitable for this algorithm and described an efficient algorithm to extract solution witnesses from true, respectively false, QBFs. We have implemented the solving and certification techniques in a tool called QuAbS. Several evaluations on different synthesis benchmarks and translated standard benchmark sets showed competitive performance. In subsequent work, we have generalized the algorithm to exploit independent quantifiers in non-prenex formulas [39] for parallel solving.

References

1. Ansótegui, C., Gomes, C.P., Selman, B.: The achilles' heel of QBF. In: Proceedings of AAAI. pp. 275–281. AAAI Press / The MIT Press (2005)
2. Balabanov, V., Jiang, J.R.: Unified QBF certification and its applications. Formal Methods in System Design 41(1), 45–65 (2012)
3. Benedetti, M., Mangassarian, H.: QBF-based formal verification: Experience and perspectives. JSAT 5(1-4), 133–191 (2008)
4. Beyersdorff, O., Bonacina, I., Chew, L.: Lower bounds: from circuits to QBF proof systems. Electronic Colloquium on Computational Complexity (ECCC) 22, 133 (2015)
5. Beyersdorff, O., Chew, L., Janota, M.: Proof complexity of resolution-based QBF calculi. Electronic Colloquium on Computational Complexity (ECCC) 21, 120 (2014)
6. Biere, A.: PicoSAT essentials. JSAT 4(2-4), 75–97 (2008)
7. Biere, A., Cimatti, A., Clarke, E.M., Zhu, Y.: Symbolic model checking without BDDs. In: Proceedings of TACAS. LNCS, vol. 1579, pp. 193–207. Springer (1999)
8. Biere, A., Lonsing, F., Seidl, M.: Blocked clause elimination for QBF. In: Automated Deduction - CADE-23 - 23rd International Conference on Automated Deduction, Wroclaw, Poland, July 31 - August 5, 2011. Proceedings. LNCS, vol. 6803, pp. 101–115. Springer (2011)
9. Bloem, R., Egly, U., Klampfl, P., Könighofer, R., Lonsing, F.: SAT-based methods for circuit synthesis. In: Proceedings of FMCAD. pp. 31–34. IEEE (2014)
10. Bloem, R., Könighofer, R., Seidl, M.: SAT-based synthesis methods for safety specs. In: Proceedings of VMCAI. LNCS, vol. 8318, pp. 1–20. Springer (2014)

11. Bohy, A., Bruyère, V., Filiot, E., Jin, N., Raskin, J.: Acacia+, a tool for LTL synthesis. In: Proceedings of CAV. LNCS, vol. 7358, pp. 652–657. Springer (2012)
12. Bradley, A.R.: SAT-based model checking without unrolling. In: Proceedings of VMCAI. LNCS, vol. 6538, pp. 70–87. Springer (2011)
13. Dershowitz, N., Hanna, Z., Katz, J.: Bounded model checking with QBF. In: Proceedings of SAT. LNCS, vol. 3569, pp. 408–414. Springer (2005)
14. Eén, N., Sörensson, N.: An extensible SAT-solver. In: Proceedings of SAT. LNCS, vol. 2919, pp. 502–518. Springer (2003)
15. Egly, U., Seidl, M., Woltran, S.: A solver for QBFs in negation normal form. Constraints 14(1), 38–79 (2009)
16. Finkbeiner, B.: Bounded synthesis for Petri games. In: Proceedings of Correct System Design. LNCS, vol. 9360, pp. 223–237. Springer (2015)
17. Finkbeiner, B., Gieseking, M., Olderog, E.: Adam: Causality-based synthesis of distributed systems. In: Proceedings of CAV. LNCS, vol. 9206, pp. 433–439. Springer (2015)
18. Finkbeiner, B., Schewe, S.: Bounded synthesis. STTT 15(5-6), 519–539 (2013)
19. Finkbeiner, B., Tentrup, L.: Detecting unrealizable specifications of distributed systems. In: Proceedings of TACAS. LNCS, vol. 8413, pp. 78–92. Springer (2014)
20. Finkbeiner, B., Tentrup, L.: Fast DQBF refutation. In: Proceedings of SAT. LNCS, vol. 8561, pp. 243–251. Springer (2014)
21. Goultiaeva, A., Bacchus, F.: Exploiting QBF duality on a circuit representation. In: Proceedings of AAAI. AAAI Press (2010)
22. Goultiaeva, A., Iverson, V., Bacchus, F.: Beyond CNF: A circuit-based QBF solver. In: Proceedings of SAT. LNCS, vol. 5584, pp. 412–426. Springer (2009)
23. Goultiaeva, A., Seidl, M., Biere, A.: Bridging the gap between dual propagation and CNF-based QBF solving. In: Proceedings of DATE. pp. 811–814. EDA Consortium San Jose, CA, USA / ACM DL (2013)
24. Heule, M., Seidl, M., Biere, A.: Efficient extraction of Skolem functions from QRAT proofs. In: Proceedings of FMCAD. pp. 107–114. IEEE (2014)
25. Jacobs, S., Bloem, R., Brenguier, R., Ehlers, R., Hell, T., Könighofer, R., Pérez, G.A., Raskin, J., Ryzhyk, L., Sankur, O., Seidl, M., Tentrup, L., Walker, A.: The first reactive synthesis competition (SYNTCOMP 2014). CoRR abs/1506.08726 (2015)
26. Janota, M., Klieber, W., Marques-Silva, J., Clarke, E.M.: Solving QBF with counterexample guided refinement. In: Proceedings of SAT. LNCS, vol. 7317, pp. 114–128. Springer (2012)
27. Janota, M., Marques-Silva, J.: Expansion-based QBF solving versus Q-resolution. Theor. Comput. Sci. 577, 25–42 (2015)
28. Janota, M., Marques-Silva, J.: Solving QBF by clause selection. In: Proceedings of IJCAI. pp. 325–331. AAAI Press (2015)
29. Jussila, T., Biere, A.: Compressing BMC encodings with QBF. Electr. Notes Theor. Comput. Sci. 174(3), 45–56 (2007)
30. Klieber, W., Sapra, S., Gao, S., Clarke, E.M.: A non-prenex, non-clausal QBF solver with game-state learning. In: Proceedings of SAT. LNCS, vol. 6175, pp. 128–142. Springer (2010)
31. Lonsing, F., Biere, A.: Nenofex: Expanding NNF for QBF solving. In: Proceedings of SAT. LNCS, vol. 4996, pp. 196–210. Springer (2008)
32. Miller, C., Scholl, C., Becker, B.: Proving QBF-hardness in bounded model checking for incomplete designs. In: Proceedings of MTV. pp. 23–28. IEEE Computer Society (2013)

33. Niemetz, A., Preiner, M., Biere, A.: Turbo-charging lemmas on demand with don't care reasoning. In: Proceedings of FMCAD. pp. 179–186. IEEE (2014)
34. Niemetz, A., Preiner, M., Lonsing, F., Seidl, M., Biere, A.: Resolution-based certificate extraction for QBF - (tool presentation). In: Proceedings of SAT. LNCS, vol. 7317, pp. 430–435. Springer (2012)
35. Pigorsch, F., Scholl, C.: Exploiting structure in an AIG based QBF solver. In: Proceedings of DATE. pp. 1596–1601. IEEE (2009)
36. QBF Gallery 2014: QCIR-G14: A non-prenex non-CNF format for quantified Boolean formulas `http://qbf.satisfiability.org/gallery/qcir-gallery14.pdf`
37. Rabe, M.N., Tentrup, L.: CAQE: A certifying QBF solver. In: Proceedings of FM-CAD. pp. 136–143. IEEE (2015)
38. Rabe, M.N., Wintersteiger, C.M., Kugler, H., Yordanov, B., Hamadi, Y.: Symbolic approximation of the bounded reachability probability in large markov chains. In: Proceedings of QEST. LNCS, vol. 8657, pp. 388–403. Springer (2014)
39. Tentrup, L.: Non-prenex QBF solving using abstraction, to appear at SAT 2016
40. Tu, K., Hsu, T., Jiang, J.R.: QELL: QBF reasoning with extended clause learning and levelized SAT solving. In: Proceedings of SAT. LNCS, vol. 9340, pp. 343–359. Springer (2015)
41. Zhang, L.: Solving QBF by combining conjunctive and disjunctive normal forms. In: Proceedings of AAAI. pp. 143–150. AAAI Press (2006)
42. Zhang, W.: QBF encoding of temporal properties and QBF-based verification. In: Proceedings of IJCAR. LNCS, vol. 8562, pp. 224–239. Springer (2014)